

Statistical Techniques in Robotics (16-831) Spring 2022

updated 4.19.2022

Days MW

Room IN-PERSON NSH 1305

Zoom: <https://cmu.zoom.us/j/95895820719> Passcode:16831

Time 11:50 - 1:10 PM

Lecturer Kris Kitani

TAs Alex LaGrassa, Kshitij Goel

Class Discussion and Slides <https://piazza.com/cmu/spring2022/16831/home>

Description

Probabilistic and learning techniques are now an essential part of building robots (or embedded systems) designed to operate in the real world. These systems must deal with uncertainty and adapt to changes in the environment by learning from experience (over time). Uncertainty arises from many sources: the inherent limitations in our ability to model the world, noise and perceptual limitations in sensor measurements, and the approximate nature of algorithmic solutions. Building intelligent machines requires that they adapt to their environment. Few things are more frustrating than machines that repeat the same mistake over and over again. We will explore (1) modern learning techniques that are effective at learning online, (2) reinforcement learning based techniques built to learn from expert demonstrations and (3) probabilistic inference algorithms for maintaining an estimate of itself despite uncertainty.

Prerequisites

Multivariate Calculus, Probability and Statistics

Grading

(1) Programming and theory assignments: 80% (5 total),

(2) Scribe notes: 20% (two scribe notes, groups of 2)

Grading scale: A: $\geq 90\%$; B: $\geq 80\%$; C: $\geq 70\%$; D: $\geq 60\%$; R: $< 60\%$;

There will be extra credit opportunities. Assignments will be submitted and graded (including feedback) via Gradescope.

Late Submissions

Please try your best to manage your time well. In the case you need extra time, you have 5 late days total for the semester. You can use up to 2 late days on a single assignments. Please use them wisely. Submissions beyond the allowed late days will be penalized by a deduction in points by $33.\bar{3}\%$ per day. Timely submissions allow for timely feedback on assignments.

Interaction and Discussion



We will be using Slack for interaction during online sessions. Students are encouraged to interact with the instructor and other students during class to increase class participation. Please try your best to engage with the lectures to enhance your learning process.



We will be using Piazza for offline class discussion. The system is highly catered to getting you help fast and efficiently learn from classmates, the TA, and myself. Rather than emailing questions to the teaching staff, I encourage you to post your questions on Piazza. If you have any problems or feedback for the developers, email team@piazza.com. Find our class sign-up link at: <https://piazza.com/cmu/spring2022/16831>. The passcode is 16831.

Educational Outcomes

(1) Able to design, implement and prove the regret bounds of a Multiplicative Weights Algorithm.

(2) Apply online learning concepts. Understand theoretical implication to incremental supervised learning algorithms. Understand online learning techniques from the perspective of online convex optimization.

(3) Implement and analyze the optimality of multi-armed bandit problems given various environments.

(4) Implement and understand the taxonomy of reinforcement learning methods for sequential learning problems.

(5) Implement and analyze the difference between linear programming, matrix games, quadratic programming and entropy maximization formulations of inverse reinforcement learning.

(6) Work in groups to summarize the essential concepts of each lecture in the form scribe notes. Develop the ability to convey theory and mathematical proofs in document form using consistent math notation and supplementary descriptions.

Academic Integrity

All encouraged to work together BUT you must do your own work (code and write up). If you work with someone, **you must include their name in your write up and inside any code that has been discussed**. If we find highly identical write-ups or code without proper accreditation of collaborators, we will take action according to university policies. For scribe notes, you are encouraged to find pre-existing presentation material to aid your preparation but you must not copy. If you use someone's material, you must give proper credit by including a citation where necessary.

Take Care of Yourself

Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress. All of us benefit from support during times of struggle. You are not alone. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is often helpful. If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit their website at <http://www.cmu.edu/counseling/>. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

Class Schedule

Date	Topic	Date	Topic
01-17	NO CLASS (MLK Day)	01-19	Robot Learning Problems
01-24	Online Learning (PWEA, Realizability)	01-26	PWEA (Halving, Randomize Greedy, Regret)
01-31	PWEA (WMA, RWMA)	02-02	OLC (Perceptron, Winnow)
02-07	OCO (Convexity), FTL	02-09	FTRL
02-14	OMD Duality (Conjugate, Bregmen Divergence)	02-16	OMD (OGD, ONEDG)
02-21	SVM (OMD/OSGD)	02-23	Boosting (OMD/EGD), MAB
02-28	Explore-Exploit, UCB	03-02	Thompson Sampling, EXP3, EXP4
03-07	NO CLASS (Spring Break)	03-09	NO CLASS (Spring Break)
03-14	Reinforcement Learning, MDP	03-16	Value-based Model-Based Control *
03-21	Model-Free Prediction (MC, N-step TD)	03-23	Model-free Prediction (lambda returns, TD-lambda)
03-28	Off-policy Prediction, Model-free Control	03-30	Function Approximation
04-04	Function Approximation, Policy Gradient	04-06	Policy Gradient, Hybrid (Actor-Critic)
04-11	Imitation Learning (Passive)	04-13	Imitation Learning (LP IRL)
04-18	Imitation Learning (Max-Margin IRL)	04-20	Imitation Learning (MaxEnt/Game IRL)
04-25	Applications in Autonomous Driving*	04-27	Imitation Learning (GAIL, Interactive)

*: guest lecturer

Assignment Schedule

HW	Topic	Release Date	Submission Date	Feedback by	Lead TA
1	Weighted Majority Algorithm	01-26	02-09	02-16	Alex and Kshitij
2	OCO and OSL	02-14	02-28	03-14	Kshitij
3	Multi-Armed Bandit	03-02	03-23	03-30	Alex
4	Reinforcement Learning	04-04	04-20	04-25	Kshitij
5	Imitation Learning	04-19	-	05-04	Alex

Scribe Notes Details

Use the Overleaf project template or template distributed by the teaching staff to compile your scribe notes. Each student will be assigned to a lecture as a note taker. It is recommended that you download the PDF of the lecture slides prior to class and take notes directly on the slides. Ask the instructor to upload the PDF prior to class if you do not see the files. Please see Piazza for the latest scribe assignments. If you need to change days to be the scribe, please negotiate directly with other students and report changes to the TA.

Guidelines:

- Do not copy and paste any portion of the lecture notes. Create the math equations yourself to make sure that you understand every part of the notation and every step of derivations. Same applies to graphics. Make your own version to ensure that you understand each aspect of the graphic.
- When needed you may refer to other material to enhance the scribe notes or to enhance your understanding. When you refer to outside material, always include a citation and explain precisely what part was used.
- Be as mathematical as possible. When appropriate use Theorems, Proofs, Corollaries, Definitions and Lemmas to structure your notes.
- Keep your language formal and concise. Follow the tone and style of a machine learning research paper or past scribe note examples.

Grade Breakdown (Total 100 pts + extra credit)

- Submit scribe notes within 3 days of the lecture (5 pt) but no later than 7 days (5 pt penalty per late day). TAs will instruct you on how to upload scribe notes.
- Add a review section briefly summarizing last lecture and give context for the present lecture. 1 page. (20 pts)
- Summary content of lecture plus citations. 5 to 7 pages (75 pts)
- Add an appendix that provides outside information, derivations, properties, inequalities, papers, books, online resources that would aid another student reading the scribe notes. (Extra Credit: 5 pts per page, up to 2 pages)

Extra Credit

- Student teams can write scribe notes for two additional lectures. Each set of scribe notes will be worth 5 points extra credit for a maximum of 10 points. Note that the extra credit scribe notes are only worth half compared to the two required scribe notes.